Cloud Service Matchmaking using Constraint Programming

Begüm İlke Zilci, Mathias Slawik, and Axel Küpper
Service-centric Networking
Technische Universität Berlin, Germany
{*ilke.zilci|mathias.slawik|axel.kuepper*}*@tu-berlin.de*

arXiv:1607.06658v1 [cs.DC] 22 Jul 2016

***Abstract*—Service requesters with limited technical knowledge should be able to compare services based on their quality of service (QoS) requirements in cloud service marketplaces. Existing service matching approaches focus on QoS requirements as discrete numeric values and intervals. The analysis of existing research on non-functional properties reveals two improvement opportunities: list-typed QoS properties as well as explicit handling of preferences for lower or higher property values. We develop a concept and constraint models for a service matcher which contributes to existing approaches by addressing these issues using constraint solvers. The prototype uses an API at the standardisation stage and discovers implementation challenges. This paper concludes that constraint solvers provide a valuable tool to solve the service matching problem with soft constraints and are capable of covering all QoS property types in our analysis. Our approach is to be further investigated in the application context of cloud federations.**

***Keywords*—Service matchmaking; Service matching; Cloud; Service marketplaces; Constraint Programming**

I. Introduction

Service brokering is a vivid research area where recent publications with several approaches can be found targeting various contexts such as cloud service marketplaces and cloud federations.

Cloud service marketplaces are platforms which act as a mediator between the service requesters and service providers. These platforms discover and store services in a service repository and allow service requesters to browse, select and interact with services via a *service broker* component. They provide a unified view of the cloud service descriptions to the service requester and in some cases additional functionality such as unified monitoring, billing, and enhanced single sign on services [1], [2], [3], [4].

Cloud federations are defined as inter-cloud organisations which comprise a set of autonomous and heterogeneous clouds [5]. Initial discussions on cloud federations bring up the question on how the service requester can keep control of the selected clouds. For this, it is suggested that the service requester should be able to set requirements to be fulfilled while the federation distributes the deployment on several clouds. A cloud federation should "respect end-to-end SLAs" (service level agreements) [6]. To ensure this, a service broker is essential.

One of the key actions of a service broker is *service matchmaking*, that is, returning one or more suitable cloud service offers from the *service providers* which fulfil the requirements of a *service requester*. Cloud service offers are any computing resources which are provided on application, system and infrastructure levels on the cloud; respectively SaaS, PaaS and IaaS. These vary in functionality, quality metrics, and other non-functional properties such as legal aspects. Moreover, the exact same application deployed on different cloud infrastructures are only different in non-functional aspects. In this setting, service matchmaking is a complex problem for which the scope of matching, the detail level of results, and the service descriptions vary. The distinguishing aspects of the services can be formally described in *service descriptions* using *service description languages* such as Linked-USDL[1] [7], SMI[2] [8], OWL-S[3] [9] and SDL-NG[4] [10] in ongoing research projects.

We apply the *Information Systems Research Framework* [11] as our research methodology. Our work primarily targets cloud service marketplaces context. Therefore, we first evaluate existing approaches against two essential requirements of service matchmaking on service marketplaces: the ability to handle incomplete knowledge and to take the service requester's perspective into account. Constraint-based approaches presented in the papers [12], [13] and [14] align at best with these requirements. Following this, we evaluate the approaches which support the essential requirements with respect to the cloud service descriptions. These approaches cover the most types of properties which result from our analysis of the service descriptions in general and in our service description language [10]. Further examination reveals two improvement opportunities: list-typed quality of service (QoS) properties as well as explicit handling of preferences for lower or higher property values.

The goal of this paper is to address these issues by explicitly handling the preferences for QoS parameters and by adding support for list typed QoS parameters. Our solution builds on the idea of using constraint programming to solve the service matchmaking problem. Therefore, we present constraint models and a prototype implementation using constraint solvers which also allows fuzzy service requests.

This paper is organised as follows: The next section describes the service matchmaking problem, the analysis of the types of properties in service descriptions, and evaluates existing approaches against the requirements. Section

[1]Linked Unified Service Description Language
[2]Service Measurement Index
[3]Semantic Markup for Web Services
[4]Next Generation Service Description Language Framework

3 presents the constraint models, Section 4 presents their implementation using the Java Constraint Programming API (JSR-331)[5]. Section 5 evaluates our approach and presents next steps.

II. Problem Definition

The requirements for service matchmakers can be defined on two levels: 1. the process as a whole, 2. the core matchmaking functionality.

The requirements for the process depend on the application context. The application context can be cloud service marketplaces, automated service composition and inter-cloud. Service marketplaces position the service matchmaker as an assistant to a service requester who is a person. Therefore, the service matchmakers in this group build up the request step by step, consider the priorities of the service requester and categorise the results as very good, good and satisfactory (*R1: Service Requester Priorities* and *R2: Comprehensive Results with Matching Degrees*).

The automatic service composition context requires that the optimal service is found without user interaction and that the over-constrained requests are automatically adjusted till a service description matches the service request. For the inter-cloud context, several ways to handle the application brokering are discussed: SLA-based, trigger-action and directly managed [5]. In the directly managed fashion, the service requester handles the deployment on multiple clouds and therefore no broker is involved. In the first two ways, a broker is involved but the service requester does not take the final decision. We argue that in a practical context the service requesters have to take the final decisions with assistance of the broker similar to service marketplaces, as they are most often liable for it - both legally and economically.

The core matchmaking functionality can be examined from three aspects: the scope of matching, incomplete knowledge and fuzziness, and the types of properties identified in the service descriptions.

The scope of matching covers non-functional and functional properties. Functional properties matching is implemented as pre- and post-condition matching and/or API matching, which rather targets the automated service composition context and software developers. The non-functional requirements of the service requester such as interoperability, quality metrics, and legal aspects are handled as QoS matching (*R3:QoS Matching*).

Service matchmakers must deal with missing values, since the service requester might not be sure about all the QoS constraints and the service provider might not supply information for all the QoS properties of the service (*R4: Incomplete Knowledge*). Moreover, the service requesters should be able to define their priorities and fuzzy values with *variational scope* for the constraints (*R5: Fuzziness*).

Variational scope [15] can be introduced to a service matching approach in three ways: (i)The matching of service description parameters to requirements with a certain amount of tolerance—assuming most requesters would accept a service with a value slightly different than the specified value, (ii) The parameters are specified with fuzzy terms such as "good" or "very fast", (iii)The requirements are specified with their level of importance to the requester with qualifiers such as mandatory and optional.

A. Analysis of QoS Properties in Service Descriptions

QoS Property	*App. X aaS by Provider #1*	*App. X aaS by Provider #2*	*App. X aaS by Provider #3*	*Service Request*
Version	5.5	5.6	5.6	$= 5.6?$
Response time	$< 120ms$	$< 200ms$	$< 400ms$	$< 300ms$
Storage in Free Version	0GB	15GB	20GB	$> 5GB$
Availability	$> 99.99\%$	$> 99.95\%$	$> 99.95\%$	$> 99\%$
Establishment Year	2010	2005	2012	> 2009
Pricing	per dyno-hour	per number of requests	per hour	per hour
Compatible Browsers	Explorer, Chrome, Firefox	Explorer, Chrome, Firefox, Safari	Explorer, Safari	Explorer, Firefox, Safari

TABLE I: An Example Service Matchmaking Problem

This paper analyses the service description languages/non-functional properties frameworks and parameters commonly used in SLA description languages, especially SMI [16] and CRF [6] [17] and the SDL-NG [10] developed as part of the project TRESOR[7] [18].

Table I shows examples for different types of properties in three service descriptions and a service request. The example describes the non-functional properties of a database application deployed on different cloud infrastructures by different cloud providers. The service request comprise the constraints on service properties on the right most column of Table I.

For *version*, the specification must be equal to the service request which is a numeric value, we name these *discrete numeric value* and *discrete value matching*. For *response time*, the service description guarantees an upper limit. It can be assumed that no discrete value matching will be performed and the lower limit of the request can be ignored. We will refer to these as *low-value preferred properties*. For *storage in free version*, higher values are preferred, only interval matching with a lower limit is needed and the service description guarantees a lower limit. Similarly, the upper limit of the request can be ignored. We name these *high-value preferred properties*. For the low-value preferred and high-value preferred properties *interval matching* is applied with the assumptions stated above. In addition, for some properties some service requesters prefer higher values and some lower values as in the case of *establishment year*. We will call this *requester defined preference*. The requester defines an upper or lower limit, and it will be matched to the service value or range e.g. $51 < x < 200$ or 100+. An example for this is the number of employees in a service provider profile. *Pricing* is an example for an enumeration, the service specification can only take one of the values in the predefined list. In the case of feature lists both the service request and the service

[5]Java Specification Request 331

[6]Cloud Requirements Framework

[7]Trusted Ecosystem for Standardized and Open cloud-based Resources

specification can take multiple values from the list as given in the *compatible browsers* example.

Based on the types of properties analysed above, we define the subproblems of service matching as follows: discrete value matching, feature list matching, interval matching, and discrete value matching with soft constraints (*R6: QoS Matching Data Types Coverage*).

III. Related Work

Some approaches do not take incomplete knowledge (R1) into account in the service descriptions and service requests [19], [20]. D'Mello et al. present an approach [21] which compares the services with each other without requester's constraints.

Table II examines QoS Matching data types coverage in related work. Although some subproblems are identified and addressed, existing approaches have shortcomings. All three approaches suggested in [22], [23] and [12] assume that the properties are either low-value or high-value preferred, although there are properties for which the service requester might be searching for exact values. The approach presented by Kritikos et al. [22], [13] improve the algorithm presented by Ruiz et al. [24] with the advanced categorisation of results. Moreover, they do not consider enumerations and fuzziness. None of them support feature lists leaving this subproblem out.

Research Work	*Features*					Maturity Level
Matchmaker	Discrete numeric data	Enumeration	Intervals	Fuzziness	Feature Lists	
[12]	yes	no	yes	yes	no	prototype
[14], [23]	yes	yes	yes	yes [a]	no	prototype and theory [b]
[22], [13], [25]	yes	no	yes	no	no	prototype
[26]	no	no	no	yes	no	theory

[a] false negatives for super matches
[b] fuzzy not in prototype

TABLE II: Data Types of Properties and Their Handling in Related Work

The interval matching approach suggested in [23] determines if the property in question is high-value preferred or low-value preferred based on the values that the service requester specified: If the most preferred value of the service requester is smaller than the least preferred value, than the property is assumed to be low-value preferred. Firstly, this sets the prerequisite that the service requester knows which values are better. However, this might not be the case. Secondly, the service requester specifies both an upper limit and a lower limit in all cases. If the property is low-value preferred, the service requester should not be prompted to specify a lower limit. Besides, the BV calculation would not work if the service requester specifies the same value for the most preferred and the least preferred values.

The trapezoidal fuzzy numbers approach [14] would deliver faulty results for low-value preferred properties and high-value preferred properties. However, it can be applied to the cases where the properties do not have broadly accepted tendencies.

IV. Constraint Models

This section presents the constraint models which address the subproblems defined in Section II-A. Bockmayr and Hooker [27] define element constraints as follows:

"$element(i,l,v)$ expresses that the $i-th$ variable in a list of variables $l=[x1,...,xn]$ takes the value $v, i.e., xi=v$."

We use element constraints to model the service matchmaking problem.

A. Model #1

Fig. 1 illustrates the first constraint model developed in this paper. Each row in the matrix contains the values of service specifications for the QoS property which is on the left most column. For each row, an element constraint is defined which adds the condition:

$$qvalues[i]"operator"qosdemand[i-1] \tag{1}$$

to the constraint solving problem. The operators can be adjusted from the list of available operators according to the specific purpose of the CSP utilising the model. The first row in the qvalues matrix is the array of service ids. For this reason, the element constraints begin with the second row. Note that the QoS request and service specifications are modelled as Java integer arrays, but not as variables, since the values for those are fixed and the variable the model sets as unknown is the index variable.

B. Model #2

The second constraint model takes another perspective to the service matching problem. Its main difference to the first model is that it takes properties as JSR-331 variables whose domain is an array consisting of the values from the service specifications. This means each row in the matrix is defined as a variable:

Listing 1: Discrete Value Matching Model

```
problem.variable( property1 ,qvalues[1]);
problem.variable( property2 ,qvalues[2]);
```

The resulting CSP searches for appropriate values of the property variables and the index variable. To ensure the integrity of a service description, additional constraints are needed, since a Java array cannot get a JSR-331 variable as an index and the index of a Java array cannot be tracked by the JSR-331. These constraints state that if serviceId has a certain value, the property value can have only one value from its domain, which is the value in domain[serviceId].

$$\begin{aligned} p &: serviceId = x, \\ q &: property1 = domain[x] \\ p &\equiv q \end{aligned} \tag{2}$$

The relation can be expressed with logical equation since p and q are either both true or both false to achieve a result true. With the element constraints, it was possible to do this without additional constraints, however element constraints are only available as hard constraints.

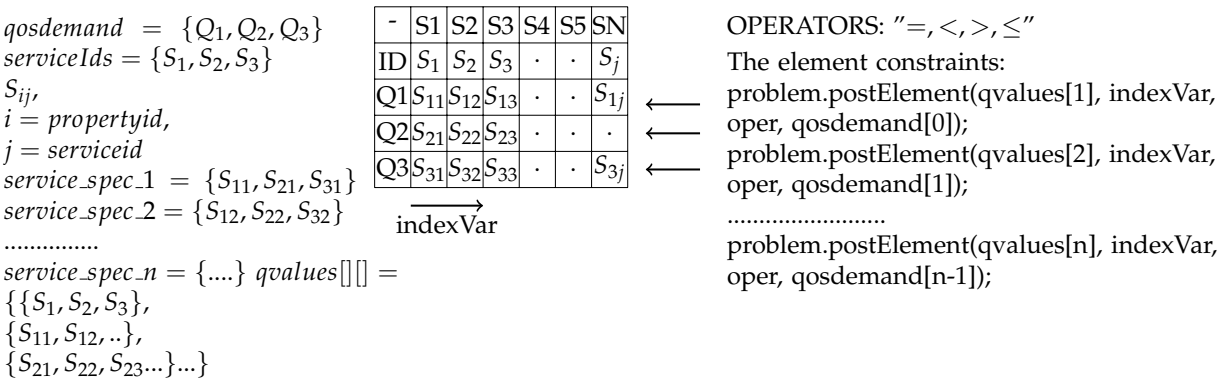

-	S1	S2	S3	S4	S5	SN
ID	S_1	S_2	S_3	·	·	S_j
Q1	S_{11}	S_{12}	S_{13}	·	·	S_{1j}
Q2	S_{21}	S_{22}	S_{23}	·	·	·
Q3	S_{31}	S_{32}	S_{33}	·	·	S_{3j}

Fig. 1: Constraint Solving Problem Model as a Matrix

V. Implementation

The prototype implementation uses JSR-331 [28] with Choco Solver [29] to implement the constraint models. Discrete value matching with hard constraints, interval matching for negative and positive tendencies and feature list matching is realised using Model #1. Discrete value matching with soft constraints is realised using Model #2. The models are described in Section IV. The implementation source can be found in our repository. [8]

Our implementation employs four methods which model the problem differently: *buildModel* for exact matching with only hard constraints, *buildModelSoftAsBool* for matching with soft Boolean constraints, *buildModelSoftDifference* for matching with soft constraints according to the difference between values of the service offer and request, and buildSimpleModelDifference which is an enhanced version of buildModelSoftDifference.

A. Discrete Value Matching with Hard Constraints

For discrete value matching with hard constraints, the implementation makes use of Model #1 which is described in Section IV. Each row in the matrix contains the discrete numeric values of service specifications. For each row there is an element constraint which adds the condition qvalues[i]=qosdemand[i-1] to the CSP. For example, indexVar=1 is in the solution set since qvalues[1][1]=2 equals to qosdemand[0].

Listing 2: Discrete Value Matching Model

```
Var indexVar = matching
.variable("serviceIndex", 0, serviceIndexMax);
for (int j = 1; j < qosdemand.length; j++) {
        matching.postElement(qvalues[j],
        indexVar, "=", qosdemand[j - 1]);
        }
```

For interval matching only the operator has to be changed: for properties with positive tendency $>=$ and for negative tendency $<=$.

[8] https://github.com/TU-Berlin-SNET/cloud-service-matcher

B. Discrete Value Matching with Soft Constraints

buildModelSoftAsBool introduces fuzziness with the third option of variational scopes described in Section II. It creates the negation of the element constraints defined in *buildModel* described in Section V-A. In contrast to hard constraints, these are not posted, instead an optimisation objective is defined using them. If the constraint is satisfied, the constraint method returns 1, if not 0. If the service specification value is not equal to the service request value the value 1 is then multiplied by the weight for the QoS parameter that the service requester specified. The violation is calculated for each element constraint and then added to the violation sum. The solver returns the service index with the minimum violation sum which is the optimisation objective.

The condition checks if the service specification value exactly matches the requirement value and if not adds up to the violation sum. In some cases, this is not enough since the requester might specify an approximate value for a requirement and the results would be still fulfilling even if they are slightly different than the requirement which is described as fuzziness with variational scope's first option above. To provide this kind of fuzziness, the optimisation objective must be the difference between the service specification value and the requirement value.

In this case, the optimisation objective can be defined as:

$$\begin{aligned} i &= qoSPropertyId \\ j &= serviceId \\ &|S_{ij} - Q_i| \end{aligned} \tag{3}$$

The coding experiments in *buildModelSoftDifference* with element constraints as soft constraints showed that if the element constraints are not posted, then the variable *serviceIndex* is not constrained, so they were not effective. For this reason, the element constraints were removed from the problem and linear constraints were added to ensure service id and service value bindings. In other words, for defining the optimisation as the difference and getting consistent results, the Model #2 was designed (see Section IV).

Getting only one solution is not suitable for the service matching problem, since it does not provide all, if there are equally optimal solutions. As a workaround, the service

matcher uses the CP solver to find all the solutions as a list and orders them according to their values for violation from minimum to maximum. This way, it can be seen if there are some solutions with the same violation value and appropriately evaluated.

C. Feature List Matching

For feature list type of constraints, a constraint solving problem per constraint must be defined.

This time, the index variable shows the elements where in the feature list QoS specification the values match with the feature list QoS constraint. We create and post a new element constraint and the default solution logger lists the values for *indexVar* and *var* which satisfy the constraint.

The matching degree is calculated based on the size of the solution set and further explained in Section V-C1. This implementation handles all the items in the required list equally.

Listing 3: Feature List Matching Example

```
String[] browsers = { "explorer", "firefox", "chrome",
 "safari", "opera" };
int[] providedBrowsers = { 1, 2, 0 };
int[] requiredBrowsers = { 0, 2, 3 };
INFO:
providedIndex[1] query[2]
providedIndex[2] query[0]
matching provided browser value:0 name:explorer
matching provided browser value:2 name:chrome
```

Listing 3 shows that the solution set has two elements: $providedBrowsers[1] = 2$ and $query = 2$, $providedBrowsers[2] = 0$ and $query = 0$. The matching degree is calculated based on the size of the solution set and further explained in Section V-C1.

1) Ranking for Feature List Matching:

Matching Degree: The Feature List Constraint contains the number codes for a list of required items. Accordingly, the Feature List QoS Specification contains the list of number codes that the service offers for that property. P is the set of provided browsers. R is the set of requested browsers. S is the set of solutions, and can be described as the intersection of provided and requested sets.

$$\begin{aligned} S = P \cap R \\ P = \emptyset \implies NOSPEC \\ S = \emptyset \implies FAIL \\ |R| > |S| \implies PARTIAL \\ |R| = |S|, |P| = |S| \implies EXACT \\ |R| = |S|, |P| > |S| \implies SUPER \end{aligned} \tag{4}$$

An example for this type of QoS property is the list of compatible browsers.

Ranking Rules: The ranking rules define how many points a service description gets according to the matching degree of its QoS specification. These are defined in the *Evaluator* classes. For example, an *ExactEvaluator* gives 2 points to the QoS specification. These rules can be changed at the corresponding Evaluator without touching other parts of the code. Table III shows the scheme for the ranking rules. At the time of writing, soft constraints calculate the violation based on the weights and the ranking for the hard constraints add scores for each matching QoS specification independently.

The final score of a service can be calculated as $(sumOfScores) - (sumOfViolations)$ as also shown below.

$$\begin{aligned} \sum_{i=1}^{j} s_i - \sum_{i=1}^{k} v_i \\ j = \text{number of hard QoS constraints} \\ k = \text{number of soft QoS constraints} \\ s_i = \text{score for the QoS Specification} \\ v_i = \text{violation for the QoS Specification} \end{aligned} \tag{5}$$

Provided	Requested	Solutions	Matching Degree	Ranking Rules
0,1,4	0,1	0,1	SUPER	3 points
0,1	0,1	0,1	EXACT	2 points
0,4	0,1	0	PARTIAL	1 point
2,3	0,1	none	FAIL	0 points
none	0,1	none	NOSPEC	0 points

TABLE III: Matching Degree Examples and Ranking Rules

VI. Evaluation

This section describes the goal-free comparison of our approach with other processes.

Our solution is designed analysing QoS properties in service descriptions, therefore it addresses *R3:QoS Matching*. Moreover, it supports the most frequent combinations of the data types in its core matching functionality addressing all sub-problems in Table II. We address *R4: Incomplete Knowledge* and *R5:Fuzziness* by allowing service requesters to define both hard and soft constraints. We provide two types of soft constraints: (i)the equality of discrete values in the specification and the request as Boolean with weights also addressing *R1: Service Requester Priorities*, (ii)the distance of the value in the specification to the specified value. *R2: Comprehensive Results with Matching Degrees* is addressed by the service matcher since it includes all evaluations of service specifications within the service descriptions which is easily accessible if needed. The implementation for the exact matching of two intervals is left for future work, since the priority was feature lists due to the TRESOR project context.

This paper contributes to the constraint-based service matching methods suggested in [13], [12] and the approaches developed in the Dino project [23], [14] by developing a better picture of the properties to be matched and diagnosing various assumptions made in the definitions of the preferences of the service requesters. It supports the view that models

Description	Design Artifact	Evaluation Method
process of matching with its inputs and outputs	method	goal-free comparison with other processes [30]
our implementation of the process	instantiation	testing

TABLE IV: Evaluation Methods

the service matching problem as an optimisation problem and that the use of constraint solvers is especially suitable for the implementation of soft constraints. It challenges the views which implicitly assume that the QoS properties are either low-value preferred or high-value preferred.

VII. Conclusion and Future Work

In this paper, we analyse the requirements for service matchmaking approaches on the process and on the core functionality levels. We identify that the application context has a defining effect on how the service requester interacts with the system and how the results are further categorised. This serves as a tool to evaluate each service matchmaker in its context. On the core functionality level, we analyse the QoS properties and identify the subproblems discrete value matching, feature list matching, interval matching, and discrete value matching with soft constraints. Our prototype implementation provides solutions for all these subproblems. We suggest that the low-value, high-value, and neutral preferences for QoS properties are explicitly stated when documenting the target properties for matchmaking functionalities.

As future work, a case study might be useful to identify additional requirements in a practical context. Moreover, we will look into the specific requirements of intercloud application brokering and extend the service matchmaker accordingly.

Acknowledgment

This work is supported by the Horizon 2020 EU funded Integrated project CYCLONE[9], grant number 644925.

References

[1] Salesforce. (2014, Oct.) App Exchange. [Online]. Available: https://appexchange.salesforce.com/

[2] Google. (2014, Oct.) Google Apps Marketplace. [Online]. Available: https://www.google.com/enterprise/marketplace/

[3] I. Amazon.com. AWS Marketplace Management Portal. [Online]. Available: https://aws.amazon.com/marketplace/management/tour

[4] D. Thatmann, M. Slawik, S. Zickau, and A. Küpper, "Towards a Federated Cloud Ecosystem: Enabling Managed Cloud Service Consumption," in *Economics of Grids, Clouds, Systems, and Services*. Springer, 2012, pp. 223–233.

[5] N. Grozev and R. Buyya, "Inter-Cloud Architectures and Application Brokering: Taxonomy and Survey," *Software: Practice and Experience*, vol. 44, no. 3, pp. 369–390, 2014.

[6] M. Assis, L. Bittencourt, and R. Tolosana-Calasanz, "Cloud Federation: Characterisation and Conceptual Model," in *Utility and Cloud Computing (UCC), 2014 IEEE/ACM 7th International Conference on*, Dec 2014, pp. 585–590.

[7] C. Pedrinaci, J. Cardoso, and T. Leidig. (2009, Jun.) linked usdl/usdl-sla. [Online]. Available: https://github.com/linked-usdl/usdl-sla

[8] C. M. U. CSMIC. Cloud Services Measures for Global Use: The Service Measurement Index (SMI). [Online]. Available: http://csmic.org/resources/

[9] D. Martin, M. Burstein, J. Hobbs, O. Lassila, D. McDermott, S. McIlraith, S. Narayanan, M. Paolucci, B. Parsia, T. Payne *et al.*, "OWL-S: Semantic Markup For Web Services," *W3C member submission*, vol. 22, pp. 2007–04, 2004.

[10] M. Slawik and A. Küpper, "A Domain Specific Language and a Pertinent Business Vocabulary for Cloud Service Selection," in *Proceedings of the 11th Conference on Economics of Grids, Clouds, Systems and Services. GECON 2014*, 2014.

[9] cyclone-project.eu

[11] A. R. Hevner, S. T. March, J. Park, and S. Ram, "Design Science Research in Information Systems," *MIS quarterly*, vol. 28, no. 1, pp. 75–105, 2004.

[12] D. Mobedpour and C. Ding, "User-centered Design of a QoS-based Web Service Selection System," *Service Oriented Computing and Applications*, pp. 1–11, 2013.

[13] K. Kritikos and D. Plexousakis, "Mixed-Integer Programming for QoS-Based Web Service Matchmaking," *IEEE Trans. Serv. Comput.*, vol. 2, no. 2, pp. 122–139, Apr. 2009. [Online]. Available: http://dx.doi.org/10.1109/TSC.2009.10

[14] D. Bacciu, M. G. Buscemi, and L. Mkrtchyan, "Adaptive Fuzzy-Valued Service Selection," in *Proceedings of the 2010 ACM Symposium on Applied Computing*. ACM, 2010, pp. 2467–2471.

[15] M. C. Platenius, M. von Detten, S. Becker, W. Schäfer, and G. Engels, "A Survey of Fuzzy Service Matching Approaches In the Context of On-The-Fly Computing," in *Proceedings of the 16th International ACM Sigsoft Symposium on Component-Based Software Engineering*. ACM, 2013, pp. 143–152.

[16] J. Siegel and J. Perdue, "Cloud Services Measures for Global Use: The Service Measurement Index (SMI)," in *SRII Global Conference (SRII), 2012 Annual*. IEEE, 2012, pp. 411–415.

[17] J. Repschlaeger, R. Zarnekow, S. Wind, and T. Klaus, "Cloud Requirement Framework: Requirements and Evaluation Criteria to Adopt Cloud Solutions," in *ECIS 2012 Proceedings. Paper 42*, 2012. [Online]. Available: http://aisel.aisnet.org/ecis2012/42

[18] S. Zickau, M. Slawik, D. Thatmann, S. Uhlig, I. Denisow, and A. Küpper, "TRESOR Towards the Realization of a Trusted Cloud Ecosystem," in *Trusted Cloud Computing*, H. Krcmar, R. Reussner, and B. Rumpe, Eds. Springer International Publishing, 2014, pp. 141–157. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-12718-7_9

[19] T. Yu, Y. Zhang, and K.-J. Lin, "Efficient Algorithms for Web Services Selection with End-to-End QoS Constraints," *ACM Transactions on the Web (TWEB)*, vol. 1, no. 1, p. 6, 2007.

[20] A. Sarang Sukumar, J. Loganathan, and T. Geetha, "Clustering Web Services Based on Multi-Criteria Service Dominance Relationship Using Peano Space Filling Curve," in *International Conference on Data Science & Engineering (ICDSE), 2012*. IEEE, 2012, pp. 13–18.

[21] D. A. D'Mello, I. Kaur, N. Ram, and V. Ananthanarayana, "Semantic Web Service Selection Based On Business Offering," in *Proceedings of the Second UKSIM European Symposium on Computer Modeling and Simulation, 2008. EMS'08.* IEEE, 2008, pp. 476–481.

[22] K. Kritikos and D. Plexousakis, "Evaluation of QoS-Based Web Service Matchmaking Algorithms," in *Proceedings of the 2008 IEEE Congress on Services - Part I*, ser. SERVICES '08. Washington, DC, USA: IEEE Computer Society, 2008, pp. 567–574.

[23] A. Mukhija, A. Dingwall-Smith, and D. S. Rosenblum, "QoS-Aware Service Composition in Dino," in *Fifth European Conference on Web Services, 2007. ECOWS'07.* IEEE, 2007, pp. 3–12.

[24] A. Ruiz-Cortés, O. Martín-Díaz, A. Durán, and M. Toro, "Improving the Automatic Procurement of Web Services Using Constraint Programming," *International Journal of Cooperative Information Systems*, vol. 14, no. 04, pp. 439–467, 2005.

[25] K. Kritikos and D. Plexousakis, "Semantic QoS-Based Web Service Discovery Algorithms," in *Proceedings of the Fifth European Conference on Web Services*, ser. ECOWS '07. Washington, DC, USA: IEEE Computer Society, 2007, pp. 181–190. [Online]. Available: http://dx.doi.org/10.1109/ECOWS.2007.21

[26] P. Wang, "QoS-Aware Web Services Selection with Intuitionistic Fuzzy Set Under Consumers Vague Perception," *Expert Systems with Applications*, vol. 36, no. 3, pp. 4460–4466, 2009.

[27] A. Bockmayr and J. Hooker, "Constraint Programming," in *Handbook of Discrete Optimization*. Elsevier, 2005, pp. 559–600. [Online]. Available: web.tepper.cmu.edu/jnh/cp-hb.pdf

[28] J. Feldman. JSR 331: Constraint Programming API, Version: 1.0.0. [Online]. Available: https://jcp.org/en/jsr/detail?id=331

[29] M. Nantes. (2014, Oct.) Choco Solver. [Online]. Available: http://www.emn.fr/z-info/choco-solver/

[30] M. Scriven, *Evaluation Thesaurus*. Sage, 1991.